

Nacelle optimisation through multi-fidelity neural networks

Francisco Sánchez-Moreno¹, David G. MacManus¹, Fernando Tejero¹ and Christopher Sheaf²

¹*Centre of Propulsion and Thermal Power Engineering, School of Aerospace, Transport and Manufacturing, Cranfield University, Bedford, United Kingdom, MK43 0AL*

²*Rolls-Royce plc., P.O. box 31, Derby, United Kingdom, DE24 8BJ*

Abstract

Purpose – Aerodynamic shape optimisation is a complex problem usually governed by transonic non-linear aerodynamics, a high dimensional design space and high computational cost. Consequently, the use of a numerical simulation approach can become prohibitive for some applications. This paper proposes a computationally efficient multi-fidelity method for the optimisation of 2D axisymmetric aero-engine nacelles.

Design/methodology/approach – The nacelle optimisation approach combines a gradient-free algorithm with a multi-fidelity surrogate model. Machine learning based on artificial neural networks (ANN) is used as the modelling technique due to its ability to handle non-linear behaviour. The multi-fidelity method combines RANS and Euler CFD calculations as high- and low-fidelity respectively.

Findings – Ratios of low- and high-fidelity training samples to degrees of freedom of $n_{LF}/n_{DOFs} = 50$ and $n_{HF}/n_{DOFs} = 12.5$ provided a surrogate model with a root mean squared error less than 5% and a similar convergence to the optimal design space when compared with the equivalent CFD-in-the-loop optimisation. Similar nacelle geometries and aerodynamic flow topologies were obtained for down-selected designs with a reduction of 92% in the computational cost. This highlights the potential benefits of this multi-fidelity approach for aerodynamic optimisation within a preliminary design stage.

Originality – The application of a multi-fidelity technique based on ANN to the aerodynamic shape optimisation problem of isolated nacelles is the key novelty of this work. The multi-fidelity aspect of the method advances current practices based on single-fidelity surrogate models and enables further computational cost reductions to meet industrial design timescales. Additionally, guidelines in terms of low- and high-fidelity sample sizes relative to the number of design variables has been established.

Keywords: multi-fidelity, neural network, surrogate model, multi-objective optimisation, nacelle.

1. Introduction

Aerodynamic shape design and optimisation is a complex problem that deals with non-linear transonic flow physics and follows a simulation-driven approach (Skinner and Zare-Behtash, 2018). For some applications the high computational overhead of the optimisation process is unaffordable within industrial timescales. One reason is the use of high fidelity computational methods to achieve realistic designs. Additionally, aerodynamic shape optimisations are typically dominated by a high number of design variables in order to allow sufficient flexibility to represent a large set of geometries (Lépine *et al.*, 2001). This increases substantially the number of designs assessed throughout the optimisation process to converge to the optimal design space (Skinner and Zare-Behtash, 2018).

Therefore, surrogate modelling and multi-fidelity techniques may be a useful approach for the aerodynamic shape optimisation problem (Leifsson and Koziel, 2017). In this respect, Wang *et al.*, 2006 applied a Kriging model to optimise the impeller of a centrifugal compressor operating at transonic conditions. The surrogate model was constructed with compressible RANS data. Relative to a baseline, the optimal design reduced the wake area and intensity of secondary flows with a resulting increase in the isentropic efficiency of 2.5%. Benaouali and Kachel, 2018 used radial basis functions for the optimisation of the ONERA M6 wing shape at $M = 0.84$ and $Re = 11.7 \times 10^6$. Relative to the baseline, the optimal design showed a reduction in C_D of approximately 15.4% for the same lift coefficient and weight. Different machine learning techniques have also been used in the context of aerodynamic shape design and optimisation. In particular, Artificial Neural Networks (ANN) have been proven to have good modelling performances for non-linear aerodynamic problems (Lindhorst *et al.*, 2014) due to the non-linear nature of the ANN back-propagation training method (Shanmuganathan and Samarasinghe, 2016). Compared to a polynomial-based techniques, the ANN models exhibit increasingly better performance as the training dataset is larger (Skinner and Zare-Behtash, 2018). In this respect, Zuccolo *et al.*, 2020 applied a Kriging and an ANN model for the prediction of the aerodynamic characteristics of exhausts afterbodies for a range of Mach numbers from 0.6 to 1.4. Both models were trained using RANS CFD data. ANN achieved better modelling performance with uncertainties in afterbody drag for 2σ confidence intervals of ± 0.010 and ± 0.013 for ANN and Kriging respectively. Bouhlel *et al.*, 2020 used ANN to perform aerodynamic shape optimisations of two different aerofoils to minimise C_D in transonic condition at $M = 0.73$. The aerofoils were parameterised using 14 design variables. The models were constructed using RANS data. The method yielded similar optimal designs to those from a CFD-in-the-loop optimisation approach with a difference of 0.12% in C_D . A reduction in computational cost of more than 95% was obtained with the ANN methodology.

Further computational cost reductions and better modelling performances can be achieved by the combination of different CFD fidelities. The low-fidelity model is used to approximate the gradients of the design space, whereas the high-fidelity one corrects these predictions. In this way, both fidelities can be combined to create surrogate models with a reduced number of high-fidelity samples needed. The low-fidelity model can be selected based on: (a) a simplified physics, (b) a coarsened spatial discretization or (c) a relaxed convergence criteria. He *et al.*, 2020 used Euler and RANS CFD data to build different surrogate models to predict the C_L and C_D of the NACA0012 aerofoil at different Mach numbers and angles of attack. Co-Kriging, multi-fidelity neural networks (MFNN) and variable complexity models (VCM) based on Kriging were used. MFNN resulted as the best performing model with a reduction in the root mean squared error (RMSE) of one order of magnitude relative to the other models. Zhang *et al.*, 2021 applied the MFNN technique to aerodynamic shape optimisation problems and compared it with a single-fidelity neural network (SFNN) approach based on high-fidelity data. RANS with boundary layer resolution and wall functions were used as high- and low-fidelity CFD models. Initially, the RAE2822 aerofoil was optimised using 12 degrees of freedom to maximize the lift-drag ratio for transonic flow conditions. For the same number of high-fidelity samples, 1% increase in lift-drag ratio relative to the baseline was obtained with MFNN-based optimisation compared to the SFNN-based optimisation. Moreover, a drag minimization of the DLR-F4 wing-body configuration in transonic flow was conducted using 30 degrees of freedom. For the same number of high-fidelity samples, the drag was reduced by 14% with the MFNN-based optimisation method whereas there was only a 3.4% reduction for the SFNN-based optimisation method.

In order to reduce the fuel consumption, the next generation of aero-engines will operate with higher bypass ratios and lower specific thrusts (Daly, 2017). This typically results in larger fan diameters and consequently an increase in the nacelle size and the overall contribution to the total weight and drag of the aircraft (Dagget *et al.*, 2003). The integration and interference effects are also expected to be exacerbated. A potential way to mitigate these aspects is

through the design of more compact nacelles. The aerodynamic shape design and optimisation of nacelles is a challenging problem governed by a high dimensional design space and a transonic non-linear aerodynamics. Compact nacelles increase further the complexity level due to the need to consider several operating points in the optimisation process to achieve robust designs to off-design conditions (Schreiner *et al.*, 2020). Consequently, the compact nacelle optimisation for high bypass ratio engines results in a highly computationally demanding optimisation problem. Surrogate models have been used to reduce the computational overhead of full CFD-based aerodynamic nacelle optimisations. Only Kriging has been used as the surrogate model for this application. Fang *et al.*, 2016 developed a hybrid genetic algorithm combining RANS numerical simulations and a Kriging surrogate model for the transonic ($M = 0.80$) aerodynamic optimisation of an axisymmetric short ($L_{nac}/r_{hi} = 3.32$) nacelle in an isolated configuration. The accuracy of the model was improved in each generation through infilling of additional RANS CFD data. Relative to a baseline geometry, the cruise drag was reduced by 4.5%. A similar approach from Tejero *et al.*, 2019 applied to the MOO of a 2D axisymmetric compact nacelle ($L_{nac}/r_{hi} = 3.0$) reduced the computational cost of the optimisation by 25% compared to a CFD-based optimisation for a similar Pareto front obtained. A multi-fidelity method was applied to the design of a compact nacelle ($L_{nac}/r_{hi} = 3.1$) considering the aerodynamic integration effects with the airframe (Tejero *et al.*, 2022). The high dimensionality of this problem was reduced from 32 to 2 design variables by applying active subspaces to an Euler CFD data base. A co-Kriging surrogate model was built in a two-dimensional space with a combination of high-fidelity RANS and low-fidelity Euler CFD calculations. Relative to a non-axisymmetric nacelle design resulting from an isolated optimisation process, an optimal nacelle was obtained with a reduction in the overall aircraft fuel consumption of about 0.6%.

It is clear that the aerodynamic shape optimisation of compact nacelles is governed by different operating conditions with highly non-linear aerodynamics which increases the computational cost of the problem. Motivated by this, the aim of this paper is to develop a methodology capable of accelerating the optimisation process of isolated nacelles to acceptable timescales within initial industrial design stages and to establish a surrogate model capable of modelling this complex flow physics. The novelty of this work is the application of a multi-fidelity technique based on ANN for aerodynamic nacelle optimisations and the establishment of guidelines in terms of low- and high-fidelity sample sizes relative to the number of design variables.

2. Methodology

The methodology developed by Tejero *et al.*, 2023 for isolated nacelle design and optimisation includes the capability of both full CFD-in-the-loop and single-fidelity surrogate-based multi-point optimisations for 2D axisymmetric nacelles. With this work the method has been extended to consider multi-fidelity surrogate modelling within the process. A MFNN approach has been established. The overall computational method encompasses the geometry parameterisation through Class Shape Transformations (CST), structured mesh generation, both viscous and inviscid compressible flow solvers for CFD evaluations and a thrust-drag accounting method for nacelle drag extraction. The optimisation is driven by an evolutionary algorithm.

2.1. Nacelle geometry parameterisation

The CST geometry parameterisation was adapted by Christie *et al.*, 2019 for nacelle applications through the use of design variables as intuitive parameters (iCST). This allows a direct correlation between the design variables and the Bernstein polynomial coefficients. The 2D axisymmetric nacelle geometry is controlled by seven degrees of freedom (Figure 1). The highlight radius (r_{hi}), the trailing edge radius (r_{te}) and the nacelle length (L_{nac}) are the key parameters that define the compactness of the nacelle and therefore are typically fixed in an optimisation problem. In this work, a compact configuration given by $L_{nac}/r_{hi} = 3.1$ and $r_{te}/r_{hi} = 0.95$ is considered. Four parameters are floated throughout the optimisation process: the maximum radius of the nacelle (r_{max}) and its axial position (f_{max}), the initial forebody radius (r_{if}) and the nacelle boat-tail angle (β_{nac}) (Figure 1). The full aero-engine is constructed by combining the nacelle geometry with a representative spinner, intake (Christie *et al.*, 2019) and separate jet exhaust (Goulos *et al.*, 2016) for ultra-high bypass ratio (UHBR) applications.




Figure 1: 2D axisymmetric nacelle parameterisation. Parameters marked in red are floated in the optimisation process.

2.2. Computational method

The design method combines two different CFD fidelities to accelerate the design process. The low-fidelity CFD model is intended to provide an estimate of the gradients in the design space by simulating the inviscid and compressible flow field. A pressure-based solver is used to resolve the Euler equations. The high-fidelity CFD model is used to account for the viscous effects and to correct the low-fidelity prediction. A density-based solver is used to solve the Reynolds-Averaged Navier Stokes (RANS) equations. Turbulent closure in the RANS equations is addressed with the κ - ω SST model. Second order upwind and Green-Gauss node-based methods are used in both models for the spatial discretization and gradients computation respectively. Sutherland's law is used to evaluate the ideal gas dynamic viscosity. To assure convergence in the calculations, a reduction of five orders of magnitude in every residual and an oscillation in the total fan-cowl force smaller than 0.05% in the last 500 iterations is required.

The 2D nacelle geometry is placed in a semi-circular domain. From a domain independence study conducted by Tejero *et al.*, 2023, a radius of 80 times the nacelle maximum radius is used. This domain is discretized using a structured C-grid mesh for both CFD fidelities. In the RANS model the first layer height is adjusted to ensure a y^+ below 1 in the wall surfaces. Using Roache's method (Roache, 1994), a GCI of about 0.3% for cruise drag is obtained from a grid independence study and a mesh with approximately 70k elements is selected. For the Euler model there is no boundary layer mesh and overall the mesh has approximately 40k elements. The RANS and Euler methods model the walls using a non-slip and slip condition respectively. The rest of the boundary conditions are common for both CFD models. The free stream conditions are imposed in the farfield through the static pressure and temperature and the flight Mach number. The static pressure and mass flow at the fan-face are prescribed to match the requirements from the engine cycle. Similarly, total pressure and temperature values are specified in both the bypass and core exhaust inlets as defined by the engine cycle. The high-fidelity RANS model was validated with experimental data available for 2D axisymmetric nacelles by Robinson *et al.*, 2018).

2.3. Optimisation approaches

Compact nacelles have large surface curvature gradients which increase the aerodynamic sensitivity to off-design conditions. Consequently, multi-point optimisation approaches are needed to ensure robust designs (Tejero *et al.*, 2023). Therefore, in addition to cruise condition (Equation 1), the sensitivities to Mach number (Equation 2) and MFCR (Equation 3) are considered. The optimisation is driven by an evolutionary algorithm. The particle swarm optimiser OMOPSO is used due to its low variability in the convergence to the optimal part of the design space for transonic and non-linear aerodynamic problems (Sánchez-Moreno *et al.*, 2022). The optimisation starts with a Design of Experiments (DoE) that uses the Latin Hypercube Sampling (LHS) technique to generate 400 non-inflected designs evenly spaced in the design space. The size of this initial generation can have a substantial impact on the scatter in the Pareto front and the convergence rate of the optimisation. Previous work by Robinson *et al.*, 2017 established 400 designs to be an acceptable value for the compact nacelle design problem. Subsequently, 39 additional generations with 50 samples per generations are performed to ensure convergence to the Pareto front. The optimisation is considered to be converged when the hypervolume change in the last three generations is smaller

than 1%. From the Pareto front a design can be down selected based on a weighted criterion of the different objective functions. From all the designs with ratios $C_{D-\Delta M}/C_{D-cruise}$ and $C_{D-EoC}/C_{D-cruise}$ lower than prescribed factors K_1 and K_2 , the design with lowest cruise drag is selected (Equation 4). Two optimisation approaches have been used. In the CFD-in-the-loop optimisation method each design in every generation is evaluated with CFD at the different operating points. An alternative is a surrogate-based optimisation approach. In this method only the initial DoE is evaluated with CFD which is used to build the surrogate model. Subsequently, OMOPSO is applied to the surrogate model to find the Pareto front without any additional CFD computational cost.

$$C_{D-cruise} = \frac{D_{nac|M=0.85,MFCR=0.7}}{\frac{1}{2}\rho_{\infty}V_{\infty}^2A_{hi}} \quad (1)$$


$$C_{D-\Delta M} = \frac{D_{nac|M=0.87,MFCR=0.7}}{\frac{1}{2}\rho_{\infty}V_{\infty}^2A_{hi}} \quad (2)$$

$$C_{D-EoC} = \frac{D_{nac|M=0.85,MFCR=0.65}}{\frac{1}{2}\rho_{\infty}V_{\infty}^2A_{hi}} \quad (3)$$


$$\min C_{D-cruise} \text{ subjected to: } \begin{cases} C_{D-\Delta M} < K_1 C_{D-cruise} \\ C_{D-EoC} < K_2 C_{D-cruise} \end{cases} \quad (4)$$

2.4. MFNN method


Artificial neural network is a machine learning technique that is suitable for complex aerodynamic problems due to the ability to handle non-linear and discontinuous data and a high number of design variables (Lindhorst *et al.*, 2014). The elemental processing unit in ANN is called a neuron. These neurons are organised in different layers to create an ANN (Figure 2a). The first and the last layers are the input and output layers, respectively. In between, a variable number of hidden layers are present which depends on the complexity of the problem to model (Shanmuganathan and Samarasinghe, 2016) (Figure 2a). In the ANN model, a certain neuron k is connected to every j neuron in the previous layer (Shanmuganathan and Samarasinghe, 2016). This can be mathematically modelled with weighting coefficients w_{kj} between every pair of neurons in different layers (Figure 2b). Moreover, a constant bias b_k is added to the neuron which helps to improve the modelling capability of the ANN (Equation 5). An activation function f is applied to the full weighted input of the neuron k to limit the amplitude of the output and enhance the training process (Shanmuganathan and Samarasinghe, 2016) (Figure 2b) (Equation 6). The choice of the activation function has a notable impact on the performance of the ANN as it determines the type of predictions that the model can make (Lau and Lim, 2019). This choice is highly problem dependant and an initial assessment has to be done for each case. Some of the most used activation functions are the rectified linear unit (ReLU), the sigmoid function and the hyperbolic tangent function. The process of learning for an ANN is based on finding the combination of weights and biases that minimise a cost function. Different cost functions like mean squared error (MSE), root mean squared error (RMSE) or mean absolute error (MAE) can be used depending on the application. The relative root mean squared error (RRMSE) is used in the present work (Equation 7). A gradient descent algorithm is used to minimise the cost function. The ADAM optimiser is used in this work due to the ability to handle large datasets and number of parameters, the computational efficiency and the fast convergence (Kingma and Ba, 2015). The ANN implementation is based on TensorFlow.



a) Representative layout of a neural network.



b) Mathematical model of a neuron.



c) Architecture of the MFNN method.

Figure 2: Neural networks modelling.

$$v_k = \sum_j w_{kj} x_j + b_k \quad (5)$$

$$y_k = f(v_k) \quad (6)$$

$$RRMSE = \sqrt{\frac{1}{n} \sum_n \left(\frac{y - y_{ANN}}{y} \right)^2} \quad (7)$$

For the multi-fidelity method two different ANNs are combined (Figure 2c) (Meng and Karniadakis, 2020). Firstly, an ANN is trained using the low-fidelity data to approximate the design space gradients. In this case the Euler CFD data is used. Secondly, an additional ANN composed of two sub-networks is fitted to capture the linear and non-linear correlation between the low- and high-fidelity models (Figure 2c). This ANN is trained using the high-fidelity RANS dataset and the predicted output from the low-fidelity ANN for this dataset. A linear activation function is used for the high-fidelity linear sub-network. For this particular application the inputs to the ANN are the selected degrees of freedom: $f_{\max}$, $r_{\max}$, r_{if} and β_{nac} . The output is the RANS-predicted nacelle drag D_{nac} (Figure 2c). A different multi-fidelity neural network model is created for each operating condition considered in the optimisation process (cruise, increased Mach number, end-of-cruise) to maximize the models performance.

Special care is needed to avoid overfitting with ANN models. Overfitting can happen for different reasons: the dataset is not large enough, the architecture of the ANN is oversized in terms of hidden layers and neurons or the training process was composed of too many back-propagation iterations (Ying, 2019). Overfitted models are not able to generalize to datasets other than the training one. For this reason, an independent dataset is used to evaluate the performance of the models and ensure generalization. Additionally, L2-norm regularization factors have been included in the MFNN technique to avoid overfitting and improve the model generalization (Shanmuganathan and Samarasinghe, 2016).

3. Verification of the MFNN method

The MFNN implementation has been verified for different benchmark cases widely used in the literature for multi-fidelity surrogate modelling. Section 3.1 shows the verification using the Forrester function (Forrester *et al.*, 2007) which is representative of a one-dimensional analytical test case. However, the correlation between the high- and the low- fidelity functions in Forrester is linear. Given the non-linear nature of the nacelle aerodynamic optimisation problem, a modified version of Forrester with non-linear correlation (called Forrester-Sanchez) have been derived and used to verify the MFNN implementation. Although not shown in this paper, the MFNN implementation was also verified for multi-dimensional functions with linear and non-linear correlation.

Considering that a random initialization is used for the weight and biases parameters in the training process, it is a common practice to perform a set of independent evaluations (N_r) and use the mean (Equation 9) and standard deviation (Equation 10) of the relative L2 error (Equation 8) to quantify the potential variability of the modelling performance. Based on previous publications (Meng and Karniadakis, 2020), a total independent evaluations from 10 to 15 are enough to capture the variability in the prediction for MFNN-based models. This has been verified in this work by a sensitivity study on N_r varying from 15 to 100 and conducted for both benchmark cases in Section 3.1 and 3.2. It has been found that the model uncertainty for a 2σ confidence level computed as $\overline{L_{2error}} \pm 2\sigma(L_{2error})$ changes by only 0.8% when increasing N_r from 15 to 100. Therefore, $N_r = 15$ has been considered in this work.

$$L_{2error} = \sqrt{\frac{\sum_n (y - y_{MFNN})^2}{\sum_n y^2}} \quad (8)$$

$$\overline{L_{2error}} = \frac{1}{N_r} \sum_{i=1}^{N_r} L_{2error_i} \quad (9)$$

$$\sigma(L_{2error}) = \sqrt{\frac{\sum_{i=1}^{N_r} (L_{2error_i} - \overline{L_{2error}})^2}{N_r}} \quad (10)$$

3.1. One-dimensional and linear correlation

The Forrester function (Equation 11 - 12) is used to verify the MFNN implementation for a one-dimensional case with linear correlation between the low- and high-fidelity datasets. For the high-fidelity level, 4 training points located at $x_{HF} = \{0, 0.4, 0.8, 1\}$ are considered. The low-fidelity training dataset is composed of 11 data points equispaced in the range from 0 to 1. This is based on a sensitivity study performed on the number of low-fidelity data points (n_{LF}) where n_{LF} has been varied from 5 to 13. The MFNN prediction performance degrades as n_{LF} reduces from 11 (Table 1). For $n_{LF} = 13$ the uncertainty reaches a plateau. The network architecture and hyperparameters have been optimized using a full factorial sampling technique to find the best performing model based on the RRMSE metric (Equation 7, Section 2.4). For the low-fidelity network, 6 hidden layers with 32 neurons per layer, hyperbolic tangent activation function and L2 regularization factor of 10^{-5} have been used. The high-fidelity network is split into the linear and non-linear sub-networks. The linear sub-network is composed of 1 hidden layer with 16 neurons and no activation function and regularization is considered. The non-linear sub-network is built using 2 hidden layers with 16 neurons per layer, hyperbolic tangent activation function and L2 regularization factor of 10^{-6} . This MFNN model provides a good agreement between the prediction and the analytical high-fidelity function (Equation 11) (Figure 3a). After 15 independent evaluations a mean L_2 relative error of 2.9% and a standard deviation of 0.9% are obtained. This equates to an uncertainty of 4.7% for a 2σ confidence level (Table 1) which is

below the 5-10% threshold commonly used for low-order model uncertainties in preliminary design stages (Heidebrecht and MacManus, 2019).

$$y_{HF}(x) = (6x - 2)^2 \sin(12x - 4) \quad (11)$$

$$y_{LF}(x) = 0.5 y_{HF}(x) + 10(x - 0.5) - 5 \quad (12)$$

n _{LF}	$\overline{L_{2error}}$	$\sigma(L_{2error})$	Uncertainty: $\overline{L_{2error}} \pm 2\sigma(L_{2error})$
5	83.0%	28.6%	140.2%
7	21.9%	3.7%	29.4%
9	13.5%	1.5%	16.5%
11	2.9%	0.9%	4.7%
13	2.7%	1.0%	4.8%

Table 1: Mean and standard deviation of the relative L2 error and 2σ uncertainty for the one-dimensional and linear correlation benchmark case for $N_r = 15$ and different number of low-fidelity training points.

3.2. One-dimensional and non-linear correlation

The Forrester-Sanchez function is a modified version of the Forrester function and is used to verify the MFNN implementation for the one-dimensional case with non-linear correlation (Equation 13 - 14). The high-fidelity training data points are increased to 6 which are equally spaced from 0 to 1. The same low-fidelity training points are used as in the linear correlation case with 11 points equi-spaced from 0 to 1. A MFNN model has been tuned to predict the high-fidelity analytical function. The low-fidelity network comprises 6 hidden layers with 32 neurons per layer, hyperbolic tangent function and L2 regularization factor of 10^{-6} . The high-fidelity linear sub-network is built with 1 hidden layer and 16 neurons. The high-fidelity sub-networks has been optimized for 3 hidden layers with 32 neurons per layer, hyperbolic tangent activation function and 10^{-6} as L2 regularization factor. The surrogate model is able to predict the global minimum at around $x = 0.75$ even when the low-fidelity dataset has a local minimum close to that position (Figure 3b). After 15 independent evaluations a mean L_2 relative error of 3.2% and a standard deviation of 1.4% are obtained. This equates to an uncertainty of 6.0% with a 2σ confidence level.

$$y_{HF}(x) = (6x - 2)\sin(12x - 4) + 5 \quad (13)$$

$$y_{LF}(x) = y_{HF}(x)^{2/3} + (x + 1)^2 - 2 \quad (14)$$




Figure 3: MFNN prediction for the one-dimensional benchmark function.

4. Results and analysis

Having demonstrated that the MFNN is able to model non-linear problems (Section 3), the same approach is applied for the aerodynamic design and optimisation of civil aero-engine nacelles. This is a complex design problem due to the large number of design variables, the nonlinearity of transonic flow aerodynamics and the high computational cost of the overall process.

4.1. MFNN hyperparameters sensitivity analysis

The optimal hyperparameters of an ANN model are highly problem dependant. The tuning process to find the best model performance requires a large number of combinations of the different parameters which can be computationally expensive. Moreover, as the MFNN technique involves different interconnected ANN (Section 2.4) there are more hyperparameters to combine and the cost to find the best model grows exponentially. This is a challenge of the method when it is first deployed on a new design problem. It can be addressed through a sensitivity analysis of the impact of different hyperparameters on the modelling performance and can help to narrow down the bounds used in the optimisation of the surrogate model.

It is not feasible to conduct the sensitivity analysis of the different combinations of parameters for the MFNN technique as it involves 9 hyperparameters to tune. In particular, 4 parameters from both the low-fidelity network and high-fidelity non-linear sub-network (activation function, number of hidden layers, number of neurons per layer and regularization factors) and 1 parameter from the high-fidelity linear sub-network (number of neurons for one layer). Therefore, just the 4 hyperparameters of the low-fidelity network are considered in the sensitivity analysis and the guidelines are extended to the corresponding parameters of the high-fidelity network. Three activation functions have been assessed: ReLu (Equation 15), sigmoid (Equation 16) and hyperbolic tangent (Equation 17). For each activation function different networks structures have been built. The number of hidden layers and neurons per layer varied from 2 to 8 and 10 to 360 respectively. Additionally, L2 norm regularization factors ranging from 10^{-3} to 10^{-6} were considered. The LHS sampling technique was used to build a training and a validation dataset with ratios of number of samples over nacelle design variables n_{LF}/n_{DOFs} of 100 and 50 respectively. The RRMSE evaluated in the validation dataset was used to assess the model predicting performance.

$$y_{ReLU}(v_k) = \max(0, v_k) \quad (15)$$

$$y_{sigmoid}(v_k) = \frac{1}{1+e^{-v_k}} \quad (16)$$

$$y_{tanh}(v_k) = \tanh(v_k) \quad (17)$$

For this design problem, the ReLu activation function outperforms sigmoid and hyperbolic tangent. For example, for the cruise operating condition, while ReLu provides models with RRMSE lower than 5%, hyperbolic tangent and sigmoid activation functions are not able to produce models with RRMSE below 30% and 50% respectively. In terms of the neural network architecture complexity, there is always a trade-off between good modelling performances and overfitting. More neurons and hidden layers tend to improve the model accuracy but may cause model overfitting (Figure 4). Low L2 regularization factors help to avoid overfitting. Based on this and on the motivation to narrow down the bounds to reduce the computational cost, a decision has been made for the hyperparameters training space to be exploited in subsequent models (Figure 4). Overall, the best modelling performances are found for a number of neurons per hidden layers between 160 and 320 (Figure 4). The number of hidden layers is kept between 2 to 4. This avoids the risk to overfitting by making the model more complex and reduces the training cost. Regularization factors from 10^{-6} to 10^{-5} are used (Figure 4). The same conclusions are drawn for the three operating conditions considered in this work (Section 2.3) (Figure 4). This highlights the low dependence of any specific hyperparameter of the neural network on the different flow physics to model. Additionally, the same sensitivity analysis has been repeated with a reduced training dataset by a factor of 4 ($n_{LF}/n_{DOFs} = 12.5$) and similar guidelines have been derived.




Figure 4: Hyperparameter sensitivity analysis of the low-fidelity network for different operating conditions. ReLu activation function. Training dataset of $n_{LF}/n_{DOFs} = 100$.

4.2. MFNN optimisation

For this study, a range of CFD datasets have been generated using the LHS technique for both the low- and high-fidelity CFD models (Section 2.2). In particular, 4 low-fidelity Euler datasets with ratios $n_{LF}/n_{DOFs} = 12.5, 25, 50$ and 100 and 5 high-fidelity RANS datasets of $n_{HF}/n_{DOFs} = 6.25, 12.5, 25, 50$ and 100 have been compiled. This rich set of data allow to assess the performance of MFNN modelling and convergence to the optimal design space by combining different ratios of n_{LF}/n_{DOFs} and n_{HF}/n_{DOFs} (Section 4.2.1). In addition to these 9 training datasets, 2 validation datasets of $n_{LF}/n_{DOFs} = n_{HF}/n_{DOFs} = 50$ have been generated.

In order to evaluate the performance of the different MFNN models on the convergence to an optimal design space, a computationally expensive RANS CFD-in-the-loop optimisation was carried out and established as reference to compare with the surrogate-based optimisations. This requires approximately 7000 2D RANS CFD calculations to achieve optimisation convergence based on the hypervolume parameter (Section 2.3). The optimal design space is represented by the non-dominated designs present in the Pareto front (Figure 5). The CFD-in-the-loop optimisation leads to designs with f_{max} and r_{max} between 20% and 40% of the normalised design space (Figure 5). Similarly, r_{if} and β_{nac} for the non-dominated designs are located below 20% of the normalised bounds (Figure 5). In particular, β_{nac} is essentially located at the lower bound of the design space. The lower bound for β_{nac} considered in this work is constrained to avoid geometry intersection between the nacelle trailing edge and the separate jet exhaust.





Figure 5: Optimal design space obtained for the CFD-in-the-loop optimisation.


4.2.1. Impact of the number of samples on the optimal design space

In multi-fidelity modelling the n_{LF}/n_{DOFs} and n_{HF}/n_{DOFs} ratios can have a substantial impact on the performance of the surrogate model and the associated computational cost. For a new problem, this is a challenge that has to be assessed as the first step through a sensitivity study. In this work, changes in the n_{LF}/n_{DOFs} and n_{HF}/n_{DOFs} ratios that span from 6.25 to 100 have been considered (Figure 6a). A MFNN surrogate model was generated for each combination of n_{LF}/n_{DOFs} and n_{HF}/n_{DOFs} . There is a trade-off between the model prediction accuracy and the associated computational overhead. As the number of samples to train the model becomes larger, the RRMSE of the model on the validation dataset decreases but the computational cost increases (Figure 6a). A common error threshold for preliminary design stages is 5%. As low-order models are typically used for preliminary design, a RRMSE up to 5% is considered acceptable in this work. This requirement is met with the lowest computational cost for $n_{LF}/n_{DOFs} = 50$ and $n_{HF}/n_{DOFs} = 12.5$ (Figure 6a).

For each combination of the multi-fidelity ratios, the trained MFNN model was used to perform a surrogate-based multi-point optimisation (Figure 6b). The resulting optimal design spaces are compared with the CFD-in-the-loop optimisation (Figure 5) to understand the impact of the ratios on the convergence to the optimal design space. For both $n_{LF}/n_{DOFs} = 12.5$ and $n_{HF}/n_{DOFs} = 12.5$ the optimisation leads to a substantially different part of the design space compared to the CFD-based reference with higher variability in the geometrical parameters (Figure 6b). This is explained by the limited predicting accuracy of the surrogate models caused by a sparse training dataset. An increase in n_{HF}/n_{DOFs} to over 25 corrects the deviations in the optimal design space at the expense of a considerable increase in the computational cost required to build the models. For example, β_{nac} is located from 0% to 30% of the investigated design space and r_{if} is not trapped in the lower bound. On the other hand, $n_{LF}/n_{DOFs} = 50$ is enough to approximate the optimal design space regardless of the high-fidelity sample size (Figure 6b). At this condition there is not a substantial benefit for $n_{HF}/n_{DOFs} > 12.5$. Therefore $n_{LF}/n_{DOFs} = 50$ and $n_{HF}/n_{DOFs} = 12.5$ is selected as the combination with the lowest computational cost that provides both acceptable surrogate model accuracy (RRMSE < 5%) and a similar convergence to the optimal design space as the equivalent CFD-in-the-loop optimisation. This surrogate-based optimisation approach provides a reduction in the overall computational cost of 92% relative to the RANS CFD-in-the-loop method.



a) Impact of the number of low- and high-fidelity samples on the computational cost and RRMSE of the surrogate model. Black cross: optimum trade-off combination between cost and model error.



- b) Optimal design spaces obtained with the surrogate-based optimisation approach for each combination of low- and high-fidelity samples. Black cross: selected combination based on a trade-off between the model accuracy, the optimal design space and the overall computational cost.

Figure 6: Computational cost, model RRMSE and optimal design space for the different combinations of low- and high-fidelity samples.

An exploration of the design space to changes in the geometrical parameters using the RANS CFD model allows an understanding of the variability on the optimal design space for the different sample ratios (Figure 7). For this purpose, a baseline design selected from the CFD-based optimal design space was perturbed for $r_{\max}$, $f_{\max}$, β_{nac} and r_{if} . The changes were applied for each design variable independently. Perturbations of 40% of the amplitude of the investigated design space were considered (Figure 7). The impacts on the nacelle drag at cruise, increased Mach number and end-of-cruise conditions were quantified relative to the baseline design. The dominant parameters are $r_{\max}$ and $f_{\max}$ with changes in nacelle drag relative to the baseline of about 90% and 80% respectively (Figure 7a, Figure 7b). Conversely, β_{nac} and r_{if} are less influential on the nacelle drag at the investigated conditions with a maximum change of 15% relative to the baseline (Figure 7c, Figure 7d). This explains the higher variability in the optimal design space for β_{nac} that results from the surrogate-based optimisations with different sample ratios (Figure 6b). The surrogate models were able to capture better the sensitivity of r_{if} to the changes in drag and therefore the optimal design space of this parameter is well defined (Figure 6b).




Figure 7: Sensitivity of the design space to changes in the design variables.

4.2.2. Down selected designs

A 3-dimensional Pareto front is obtained which is represented by a 2D projection on the $C_{D\text{-cruise}}$ and $C_{D\text{-}\Delta\text{M}}$ plane and coloured by $C_{D\text{-EoC}}$ (Figure 8). The Pareto fronts that result from the CFD-in-the-loop optimisation (Figure 8a) and the selected surrogate-based optimisation method ($n_{\text{LF}}/n_{\text{DOFs}} = 50$ and $n_{\text{HF}}/n_{\text{DOFs}} = 12.5$) (Figure 8b) show a similar relationship between metrics. While the drag data on the CFD-based Pareto come from RANS CFD calculations, the drag on the surrogate-based Pareto is predicted with the MFNN model. From each Pareto front a design was selected using the same selection criteria (Section 2.3, Equation 4): Design D1 from the CFD-based optimisation and Design D2 from the MFNN-based optimisation (Figure 8). Similar drag reductions for each operating condition relative to a baseline design ($(C_{D\text{-selected}} - C_{D\text{-baseline}})/C_{D\text{-baseline}}$) were obtained for both the CFD-based and MFNN-based designs. A difference of just 1% in the drag reduction is obtained between both optimisation methods.




Figure 8: Pareto fronts and down selected designs from the optimisations using the CFD-in-the-loop and the surrogate-based approaches.

To quantify the accuracy of the prediction of the MFNN model in the optimal design space, a cross-validation study shows the agreement between the surrogate model predictions and the corresponding CFD results for the designs contained in the Pareto front obtained on the MFNN-based optimisation (Figure 9). For each model at every flight condition, most of the designs are within the $\pm 5\%$ bounds. The RRMSE of the cruise, increased Mach number and end-of-cruise models is reduced from approximately 5% in the full validation dataset (Figure 6a) to 1.40%, 2.5% and 1.6% respectively in the optimal design space. This highlights the reliability of the MFNN model in the selection of the optimal design.




Figure 9: Surrogate model performance in the optimal design space using the designs in the Pareto front obtained in the MFNN-based optimisation.

Very good agreement of the design variables has been achieved for the two selected designs (D1 and D2). The only noticeable differences are for $r_{\max}$ and r_{if} with small reductions of $0.012r_{\text{hi}}$ and $0.003r_{\text{hi}}$ for the MFNN-based selected design relative to the CFD-in-the-loop optimal geometry. Similar flow physics are obtained for both nacelle designs using a RANS CFD evaluation (Figure 10). A slight overprediction of approximately 0.1 in the peak isentropic Mach number is apparent for the MFNN-based design which is caused by the $0.003r_{\text{hi}}$ reduction in r_{if} . At the cruise condition, the same flow topology is exhibited by both geometries. The flow acceleration close to the nacelle crest is captured in the aerodynamics of both designs with an axial offset of $0.05L_{\text{nac}}$ approximately (Figure 10). The same flow topology is also captured for the increased Mach number operating condition with a forward movement of the shock of $0.05L_{\text{nac}}$ approximately for the surrogate-based selected design (Figure 10). For the end-of-cruise operating condition, similar shock location and pre-shock Mach number are obtained with both designs. There is a slight underprediction in the flow reacceleration after the shock of approximately $0.05M_{\text{is}}$ for the surrogate-based design (Figure 10).

Overall, similar geometry with similar flow physics and drag performance benefits are obtained with both optimisation approaches with a significant reduction of 92% in the computational cost for the proposed MFNN

optimisation method relative to the CFD-in-the-loop approach. This outcome has a substantial impact in the industry sector, which is mostly limited by short timescales. The speed-up achieved in the method allows to explore a wider design space and understand the sensitivities. Additionally, a key element in nacelle design for industrial standard is the use of geometrical hard-points to accommodate the different engine components. Therefore, this method can faster assess the impact of the different hard-points on the optimal region of the design space and understand the trade-off between aerodynamic performance and hard-points.

The implementation of the method is completely compatible with the optimisation problem of non-axisymmetric compact 3D nacelles. The challenge is on the wider design space with increased number of design variables and the required amount of low- and high-fidelity training data points for the MFNN. Although in this 2D optimisation work $n_{LF}/n_{DOFs} = 50$ and $n_{HF}/n_{DOFs} = 12.5$ has been selected as the optimal combination, a similar study is required for the 3D problem given the higher dimensionality and the more complex correlation between design variables. Alternatively, dimensionality reduction techniques can help to identify the dominant design variables upon which the MFNN model can be constructed and in this way reduce the overall computational cost of the problem. In general, this approach can be applied to any aerodynamic design and optimisation problem led by overall integral metrics like drag and lift coefficients or isentropic efficiencies.




Figure 10: Mach number distributions for the down selected designs using the CFD-in-the-loop and surrogate-based optimisation approaches.

5. Conclusions

This work has presented a method to perform multi-point aerodynamic shape optimisations for transonic non-linear applications using a multi-fidelity technique based on neural networks. The method has been applied to the optimisation of 2D axisymmetric compact nacelles for ultra-high bypass ratio aero-engines. The impact of the low- and high-fidelity training samples to degrees of freedom ratios on the surrogate model predicting performance, convergence to the optimal design space and associated computational cost has been assessed. As a result, $n_{LF}/n_{DOFs} = 50$ and $n_{HF}/n_{DOFs} = 12.5$ has been selected as the combination with the lowest computational cost that allows

surrogate models with a RRMSE $< 5\%$ in the validation dataset and a similar convergence to the optimal design space compared with the equivalent CFD-in-the-loop optimisation. The comparison of the optimal designs obtained from both optimisation approaches using the same down selection criteria shows similar nacelle geometric parameters with similar flow physics. The proposed method reduces the computational cost to perform the multi-point optimisation by 92% relative to the CFD-in-the-loop approach. This shows that the MFNN surrogate-based optimisation methodology is suitable for transonic aerodynamic shape optimisation problems within preliminary design stages.

6. Acknowledgements

This project has received funding from the Clean Sky 2 Joint Undertaking under the European Unions Horizon 2020 Research and Innovation Program under Grant Agreement No 820997.



References

- Benaouali, A. and Kachel, S. (2018), "A surrogate-based integrated framework for the aerodynamic design optimization of a subsonic wing planform shape", *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, Vol. 232 No. 5, pp.872–883. doi:10.1177/0954410017699007.
- Bouhlel, M.A. and He, S. and Martins, J.R.R.A. (2020), "Scalable gradient enhanced artificial neural networks for airfoil shape design in the subsonic and transonic regimes", *Structural and Multidisciplinary Optimization*, Vol. 61 No. 4, pp.1363–1376. doi:10.1007/s00158-020-02488-5.
- Christie, R. and Robinson, M. and Tejero, F. and MacManus, D. (2019), "The use of hybrid intuitive class shape transformation curves in aerodynamic design", *Aerospace Science and Technology*, Vol. 95.
- Daggett, D. and Brown, S. and Kawai, R. (2003), *Ultra-efficient engine diameter study*, NASA Technical Report No. 2003-212309.
- Daly, M. (2017), *Jane's Aero Engines*, IHS Jane's.
- Fang, X. and Zhang, Y. and Li, S. and Chen, H. (2016), "Transonic nacelle aerodynamic optimization based on hybrid genetic algorithm", *17th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Washington, D. C., 13-17 June*. doi:10.2514/6.2016-3833.
- Forrester, A.I.J. and Sóbester, A. and Keane, A.J. (2007), "Multi-fidelity optimization via surrogate modelling", *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, Vol. 463 No. 2088, pp.3251–3269. doi:10.1098/rspa.2007.1900.
- Goulos, I. and Stankowski, T. and Otter, J. and MacManus, D. and Grech, N. and Sheaf, C. (2016) "Aerodynamic design of separate-jet exhausts for future civil aero-engines – part I: parametric geometry definition and computational fluid dynamics approach", *Journal of Engineering for Gas Turbines and Power*, Vol. 138 No. 8, pp.1–14. doi:10.1115/1.4032649.

He, L. and Qian, W. and Zhao, T. and Wang, Q. (2020), “Multi-fidelity aerodynamic data fusion with a deep neural network modeling method”, *Entropy*, Vol. 22 No. 9. doi:10.3390/e22091022.

Heidebrecht, A. and MacManus, D.G.(2019), “Surrogate model of complex non-linear data for preliminary nacelle design”, *Aerospace Science and Technology*, Vol. 84, pp.399–411. doi:10.1016/j.ast.2018.08.020.

Kingma, D.P. and Ba, J. (2015), “Adam: a method for stochastic optimization”, *International Conference on Learning Representation, San Diego, California, 7-9 May*. doi: 10.48550/arXiv.1412.6980.

Lau, M.M. and Lim, K.H. (2019), “Review of adaptive activation function in deep neural network”, *2018 IEEE-EMBS Conference on Biomedical Engineering and Sciences, Kuching, Sarawak, Malaysia, 3-6 December*. doi:10.1109/IECBES.2018.08626714.

Lépine, J. and Guibault, F. and Trepanier, J.Y. and Pépin, F. (2001), “Optimized nonuniform rational B-spline geometrical representation for aerodynamic design of wings”, *AIAA Journal*, Vol. 39 No. 11, pp.2033–2041.

Lindhorst, K. and Haupt, M.C. and Horst, P. (2014), “Efficient surrogate modelling of nonlinear aerodynamics in aerostructural coupling schemes”, *AIAA Journal*, Vol. 52 No. 9, pp.1952–1966. doi:10.2514/1.J052725.

Leifsson, L. and Koziel, S. (2017), “Adaptive response prediction for aerodynamic shape optimization”, *International Journal of Numerical Methods for Heat and Fluid Flow*, Vol. 34 No. 5, pp.1485-1500. doi: 10.1108/EC-02-2016-0070.

Meng, X. and Karniadakis, G.E. (2020), “A composite neural network that learns from multi-fidelity data: Application to function approximation and inverse PDE problems”, *Journal of Computational Physics*, Vol. 401. doi:10.1016/j.jcp.2019.109020.

Roache, P. (1994), “A method for uniform reporting of grid refinement studies”, *Journal of Fluids Engineering*, Vol. 116 No. 3, pp.405–413.

Robinson, M.H. and MacManus, D.G. and Heidebrecht, A. and Grech, N. (2017), “An Optimisation Method for Nacelle Design”, *55th AIAA Aerospace Sciences Meeting, Grapevine, Texas, 9-13 January*.

Robinson, M. and MacManus, D. and Sheaf, C. (2018), “Aspects of aero-engine nacelle Drag”, *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, Vol. 233 No. 5, pp.1667–1682.

Sánchez-Moreno, F. and MacManus, D. and Tejero, F. and Matesanz-Garcia, J. and Sheaf, C. (2022), “Robustness of optimisation algorithms for transonic aerodynamic design”, *9th European Conference for Aeronautics and Space Sciences, Lille, France, 27 June- 1 July*.

Schreiner, B.D.J and Tejero, F. and MacManus, D.G. and Sheaf, C. (2020), “Robust aerodynamic design of nacelles for future civil aero-engines”, *Proceedings of ASME Turbo Expo 2020: Power for Land, Sea and Air, virtual, 21-25 September*, GT2020-14470, doi:10.1115/GT2020-14470.

Shanmuganathan, S. and Samarasinghe, S. (2016), *Artificial Neural Network Modelling*, Springer.

Skinner, S. and Zare-Behtash, H. (2018), “State-of-the-art in aerodynamic shape optimisation methods”, *Applied Soft Computing*, Vol. 62, pp.933–962. doi:10.1016/j.asoc.2017.09.030.

Tejero, F. and MacManus, D. and Hueso-Rebassa, J. and Sanchez-Moreno, F. and Goulos, I. and Sheaf, C. (2022), “Aerodynamic optimisation of civil aero-engine nacelles by dimensionality reduction and multi-fidelity techniques”, *International Journal of Numerical Methods for Heat and Fluid Flow*, Vol. 33 No. 4, pp.1336-1353. doi:10.1108/HFF-06-2022-0368.

Tejero, F. and MacManus D.G and Matesanz-Garcia, J. and Swarthout, A. and Sheaf, C. (2023), "Towards the design and optimisation of future compact aero-engine: intake/fan-cowl trade-off investigation", *International Journal of Numerical Methods for Heat and Fluid Flow*, Vol. 33 No. 4, pp.1319-1335. doi:10.1108/HFF-06-2022-0366.

Tejero, F. and MacManus, D.G. and Sheaf, C. (2019), "Surrogate-based aerodynamic optimisation of compact nacelle aero-engines", *Aerospace Science and Technology*, Vol. 93. doi:10.1016/j.ast.2019.05.059.

Wang, X.F. and Xi, G. and Wang, Z.H. (2006), "Aerodynamic optimization design of centrifugal compressor's impeller with Kriging model", *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, Vol. 220 No. 6, pp.589–597. doi:10.1243/09576509JPE201.

Ying, X. (2019), "An overview of overfitting and its solutions", *Journal of Physics: Conference Series*, Vol. 1168 No. 2. doi:10.1088/1742-6596/1168/2/022022.

Zhang, X. and Xie, F. and Ji, T. and Zhu, Z. and Zheng, Y. (2021), "Multi-fidelity deep neural network surrogate model for aerodynamic shape optimization", *Computer Methods in Applied Mechanics and Engineering*, Vol 373. doi:10.1016/j.cma.2020.113485.

Zuccolo, G. and Christie, R. and MacManus, D. and Goulos, I. and Martin, P. (2020), "Low order models for transonic afterbody aerodynamic characteristics", *AIAA Science and Technology Forum and Exposition, Orlando, Florida, 6-10 January*, pp.1–16. doi:10.2514/6.2020-1997.



DSPACE

<https://dspace.org/>

Nacelle optimisation through multi-fidelity neural networks

Sánchez-Moreno, Francisco; MacManus, David G.; Tejero, Fernando; Sheaf, Christopher T.

2024-09-04

Emerald

<https://doi.org/10.1108/HFF-12-2023-0745>

Sánchez-Moreno F, MacManus D, Tejero F, Sheaf C. (2024) Nacelle optimisation through multi-fidelity neural networks. International Journal of Numerical Methods for Heat & Fluid Flow, Volume 34, Issue 9, September 2024, pp. 3615-3634